

IN THE SPECIFICATION

Please amend the Abstract on Page 36, lines 2-27 as follows:

A method for performing a join of a first table and a second table is disclosed. In one instance, each of the plurality of first table rows is grouped into one of a plurality of first table partitions. It is determined that a partitioning column of the first table is specified in an equality condition of the join. One of the plurality of second table rows is selected, and a partition number is generated for the selected one of the plurality of second table rows. For a second table row for which a partition number was generated, a first table partition with a partition number that matches the generated partition number is identified, and a product join is used to join the one or more first table rows in the matching partition with the selected one of the second table rows. Methods, computer programs and database systems for performing a join are disclosed. The join includes a right table and a left table. The right table includes a plurality of right table rows. In one instance, each of the plurality of right table rows is grouped into one of a plurality of right table partitions. The left table includes one or more left table rows. It is determined that a partitioning column of the right table is specified in an equality condition of the join. The left table is prepared to join one or more rows of the left table with one or more rows of the right table by selecting one of the plurality of left table rows and generating a partition number for the selected one of the plurality of left table rows. For a left table row for which a partition number was generated, a right table partition with a partition number that matches the generated partition number is identified, and a product join is used to join the one or more right table rows in the matching partition with the selected one of the left table rows if one or more join conditions are satisfied. In another instance, the left and right tables are partitioned, the left table and the right table are joined on equality constraints and a relationship is identified between a partitioning expression associated with the left table and a partitioning expression associated with the right table. It is determined that each left table partition matches at least one right table partition based on the relationship between the partitioning expressions associated with the left table and

~~right table. The one or more left table rows of each matching left table partition are joined with the one or more right table rows of the at least one matching right table partition if one or more join conditions are satisfied. In another instance, the join specifies an equality constraint between each primary index column of the left table and a corresponding primary index column of the right table and inequality conditions between each partitioning column of the left table and a corresponding partitioning column of the right table. A mapping of the plurality of left table partitions to the plurality of right table partitions is determined to be a one to many relationship and at least one left table partition is determined to match at least two right table partitions based on the relationship between the partitioning expressions associated with the left table and right table. The one or more left table rows of the at least one matching left table partition are joined with the one or more right table rows of the at least two matching right table partitions if one or more join conditions are satisfied.~~

Please amend the paragraph beginning "In general, in another aspect, the invention features a database system including a massively parallel processing system . . ." on Page 7, line 23 of the Specification as follows:

[0015] In general, in another aspect, the invention features a database system including a massively parallel processing system, which includes one or more nodes, a plurality of CPUs, each of the one or more nodes providing access to one or more CPUs, a plurality of data storage facilities each of the one or more CPUs providing access to one or more data storage facilities, a process for execution on the massively parallel processing system for performing a join. The process includes identifying a join that identifies a right table and a left table. The right table includes a plurality of right table rows. Each of the plurality of right table rows is grouped into one of a plurality of right table partitions. The left table includes one or more left table rows. ~~the process~~ The process further includes determining that a partitioning column of the right table is specified in an equality condition of the join and preparing the left table for joining one

or more rows of the left table with one or more rows of the right table. Preparing the left table includes selecting one of the plurality of left table rows and generating a partition number for the selected one of the plurality of left table rows. The process further includes, for a left table row for which a partition number was generated, identifying a right table partition with a partition number that matches the generated partition number and joining, using a product join, the one or more right table rows in the matching partition with the selected one of the left table rows if one or more join conditions are satisfied.

Please amend the paragraph beginning "The performance of joins of partitioned tables may be improved . . ." on Page 14, line 15 of the Specification as follows:

[0044] The performance of joins of partitioned tables may be improved, however, when certain conditions are present. For example, performance of the join may be improved where there are equality conditions on both the primary index and partitioning columns of tables 505 and 515. For example, where tables 505 and 515 are partitioned using the same partition function and tables 505 and 515 are joined on the partitioning column, certain join efficiencies may be realized as described in co-pending United States Patent ~~Application Serial No. _____~~
~~No. 6,944,633~~, entitled Performing a ~~Rowkey Merge~~ Join in a Partitioned Database System, by ~~Paul Sinclair, Kuorong Chiang, Larry Higa and Mark Sirek~~ Lawrence H. Higa, Paul L. Sinclair, Mark William Sirek, and Kuorong Chiang, NCR Docket No. 11276, filed on December 10, 2003.

Please amend the paragraph beginning "The system then joins each left table row . . ." on Page 16, line 23 of the Specification as follows:

[0049] The system then joins each left table row with the right table. It does this by selecting each left table row in turn (block 630) and identifying ~~a the right the right~~ table partition with a partition number matching that generated for the selected left table row (block 635) and joining, using a product join, the right table rows in the matching partition with the selected left table

ATTORNEY DOCKET NO.
11196

PATENT APPLICATION
SERIAL NO. 10/748,563

row if one or more join conditions are satisfied (block 640). If all of the left table rows have been processed (block 645), the processing is complete (block 650). Otherwise, the system selects a new left table row (block 635) and repeats the processing described above.